



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/828,573	04/21/2004	Matthew A. Ahrens	03226.390001; P8399	5289
32615	7590	10/22/2007	EXAMINER	
OSHA LIANG L.L.P./SUN			DEBNATH, SUMAN	
1221 MCKINNEY, SUITE 2800				
HOUSTON, TX 77010			ART UNIT	PAPER NUMBER
			2135	
			NOTIFICATION DATE	DELIVERY MODE
			10/22/2007	ELECTRONIC

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.

Notice of the Office communication was sent electronically on above-indicated "Notification Date" to the following e-mail address(es):

lord@oshaliang.com
hernandez@oshaliang.com
DOCKETING@OSHALIANG.COM

Office Action Summary	Application No.	Applicant(s)
	10/828,573	AHRENS ET AL.
	Examiner	Art Unit
	Suman Debnath	2135

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --
Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) OR THIRTY (30) DAYS, WHICHEVER IS LONGER, FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 01 August 2007.
 2a) This action is FINAL. 2b) This action is non-final.
 3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11; 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1-26 is/are pending in the application.
 4a) Of the above claim(s) _____ is/are withdrawn from consideration.
 5) Claim(s) _____ is/are allowed.
 6) Claim(s) 1-26 is/are rejected.
 7) Claim(s) _____ is/are objected to.
 8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.
 10) The drawing(s) filed on 04/21/2004 is/are: a) accepted or b) objected to by the Examiner.
 Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
 Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).
 11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).
 a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
 3) Information Disclosure Statement(s) (PTO/SB/08)
 Paper No(s)/Mail Date 08/01/2007.

4) Interview Summary (PTO-413)
 Paper No(s)/Mail Date. _____.
 5) Notice of Informal Patent Application
 6) Other: _____.

DETAILED ACTION

1. Claims 1-26 are pending in this application.
2. Claims 1-2, 4-5, 10-11, 15, 17-18, 24-26 are presently amended.
3. The text of those sections of Title 35, U.S. Code not included in this action can be found in a prior Office Action.

Claim Rejections - 35 USC § 102

4. Claims 1-26 are rejected under 35 U.S.C. 102(b) as being anticipated by Talagala et al. (Pub. No.: US 2002/0161972 A1) (hereinafter "Talagala").
5. As to claim 1, Talagala discloses a method for storing data blocks (abstract), comprising: storing a first data block and a second data block in a storage pool (Talagala teaches this concept by providing parity/stripe group tables having entries for multiple blocks of data stored in a storage pool, -e.g. [0059], FIG. 6C, 7B and 8B); obtaining a first data block location and a second data block location ([0017], [0059]); calculating a first data block checksum for the first data block (Talagala discloses writing blocks having checksum and location in PGT (Parity Group Table) which have "segment" filed to store locations and "checksum" fields to store checksum, -e.g. [0059], FIG. 6C, 7B and 8B); calculating a second data block checksum for the second data block ([0059], FIG. 6C, 7B and 8B); and storing a first indirect block referencing the first data block and the second data block in the storage pool, wherein the first indirect block comprises a first block pointer comprising the first data block location and the first data

Art Unit: 2135

block checksum and a second block pointer comprising the second data block location and the second data block checksum ("an indirection map (e.g., block remapping table) matches virtual block address to physical block address. Block-level checksums may be provided in the indirection map" —e.g. [0059] and FIG. 6C; "each valid PGT entry also includes a back pointer to the next entry in a parity group so that the first physical segment in a parity group is linked to the second physical segment in that parity group, and the second physical segment to the third and so on, until the last physical segment contains the parity data for that parity group. The physical segment that contains the parity data is linked back to the first physical segment in the parity group, thereby creating a circular list for that parity group" —e.g. [0056], Applicant should note that a circular linked list will comprise "a first physical segment in a parity group" which comprises a first indirect block referencing the first and second data blocks and so on as claimed by applicant).

6. As to claim 2, Talagala discloses further comprising: calculating a first indirect block checksum for the first indirect block ([0059], [0061]); obtaining a first indirect block location ([0059]); and storing a second indirect block in the storage pool, wherein the second indirect block comprises the first indirect block location and the first indirect block checksum (Talagala teaches writing a block having a checksum and a location in PGT (Parity Group Table) which have a "segment" field to store a location and a "checksum" field to store a checksum and provides which contain parity/stripe group

tables having entries for multiple blocks of data stored in a storage pool, -e.g. [0059] and FIG. 6C, 7B and 8B).

7. As to claim 3, Talagala discloses further comprising: assembling the first indirect block ("when a block read is requested, the block's indirection map entry is read to find the block's physical address" -e.g. [0061], lines 8-10).

8. As to claim 4, Talagala discloses wherein assembling the first indirect block comprises: storing the first data block checksum in a checksum field in the first block pointer in the first indirect block (Talagala teaches this concept as "each block's checksum is stored in an entry in the indirection map, e.g. as part of a block remapping table entry (e.g. PGT entry) for each block" -e.g. [0058] and FIG. 6C, 7B and 8B, wherein each entry representing/pointing to a block in PGT contains a checksum in a checksum field), and storing the first data block location in the first block pointer, wherein storing the data block location comprises storing a metaslab ID (Talagala teaches this concept as "the indirection map may also include a parity group pointer for each data block that points to a next member of that parity group" -e.g. [0013], wherein "when a READ or WRITE command is received for a block(s), the appropriate PGT entry is accessed to locate the blocks in the disk drives" -e.g. [0058], "HIT (Hash Indirection Table)" and explains having "PGT (Parity Group Table)" indices to access a PGT which contains configuration information for each of the parity groups such as a segment field which indicated the physical disk location (disk and segment) of parity

Art Unit: 2135

groups –e.g. [0055], [0058], [0059] and FIG. 6B, 6C, 7A, 7B, 8A and 8B). Talagala provides an example in which Virtual address 0 corresponds to PGT index 12, which contains valid data at physical segment D1.132 wherein "this may be interpreted as Disk 1, segment 132" –e.g. [0057] and Figures 6B, 6C, 7A, 7B, 8A and 8B). As defined by Applicant, metaslabs are "contiguous regions of data" in which "the storage space in the storage pool is divided" (Specification, [0032]); Therefore, Applicant should note that these metaslabs may comprise any amount of contiguous data, such as segments or blocks into which a storage pool is divided, as described by Talagala) and an offset (Talagala teaches this concept as entries stored under the "next entry in Parity Group" field in PGT (Parity Group Table) which points to the next virtual block entry, -e.g. [0057] and FIG. 6B, 6C, 7A, 7B, 8A and 8B) and offset (Talagala teaches this concept as entries stored under the "next entry in Parity Group" field in PGT (Parity Group Table) which points to the next virtual block entry, -e.g. [0057] and FIG. 6B, 6C, 7A, 7B, 8A and 8B).

9. As to claim 5, Talagala discloses further comprising: storing a birth value in a birth field in the first block pointer ("a hashed indirection table (HIT) which maintains generational images" and explains that "the PGT index columns are now labeled version zero through version two, where version zero corresponds to the most current version and version two corresponds to the oldest version" –e.g. [0065]).

Art Unit: 2135

10. As to claim 6, Talagala discloses wherein the first indirect block is assembled using a data management unit ("Storage Controller 401" –e.g. FIG. 2 and 3, [0033]).

11. As to claim 7, Talagala discloses wherein the storage pool comprises at least one storage device ("array of storage devices 410" –e.g. FIG. 2 and 3, [0033]).

12. As to claim 8, Talagala discloses wherein the storage pool is divided into a plurality of metaslabs (Talagala teaches having different "stripes of data" within an array storage devices –e.g. [FIG. 4 and [0043] and explains that a "stripe" of data is analogous to a "parity group" –e.g. [0063], lines 10-11, wherein configuration information for each of the "parity groups" is stored in a "PGT (parity group table)" –e.g. FIG. 6C, 7B and 8B, and further explains "the indirection map may also include a parity group pointer for each data block that points to a next member of that parity group" – e.g. [0013], wherein "when a READ or WRITE command is received for a block(s), the appropriate PGT entry is accessed to locate the blocks in the disk drives" –e.g. [0058]. As defined by Applicant, metaslabs are "contiguous regions of data" in which "the storage space in the storage pool is divided" (Specification, [0032]); therefore, Applicant should note that these metaslabs may comprise any amount of contiguous data, such as segments or blocks into which a storage pool is divided, as described by Talagala).

13. As to claim 9, Talagala discloses wherein each of the plurality of metaslabs is associated with a metaslab ID (Talagala teaches this concept as each virtual block has

Art Unit: 2135

a virtual address which is used to access a "HIT (Hash Indirection Table)" which contains "PGT (Parity Group Table) indices to access a PGT which contains configuration information for each of the parity groups/metaslabs such as a segment field which indicated the physical disk location (disk and segment) of parity groups. – e.g. FIG. 6B, 6C, 7A, 7B, 8A and 8B, [0055], [0058], [0059]).

14. As to claim 10, Talagala discloses wherein the first data block location comprises the metaslab ID (Talagala teaches this concept as "the indirection map may also include a parity group pointer for each data block that points to a next member of that parity group" –e.g. [0013], wherein "when a READ or WRITE command is received for a block(s), the appropriate PGT entry is accessed to locate the blocks in the disk drives" – e.g. [0058], "HIT (Hash Indirection Table)" and explains having "PGT (Parity Group Table)" indices to access a PGT which contains configuration information for each of the parity groups such as a segment field which indicated the physical disk location (disk and segment) of parity groups –e.g. [0055], [0058], [0059] and FIG. 6B, 6C, 7A, 7B, 8A and 8B). Talagala provides an example in which Virtual address 0 corresponds to PGT index 12, which contains valid data at physical segment D1.132 wherein "this may be interpreted as Disk 1, segment 132" –e.g. [0057] and Figures 6B, 6C, 7A, 7B, 8A and 8B). As defined by Applicant, metaslabs are "contiguous regions of data" in which "the storage space in the storage pool is divided" (Specification, [0032]); Therefore, Applicant should note that these metaslabs may comprise any amount of contiguous data, such as segments or blocks into which a storage pool is divided, as described by

Art Unit: 2135

Talagala) and an offset (Talagala teaches this concept as entries stored under the "next entry in Parity Group" field in PGT (Parity Group Table) which points to the next virtual block entry, -e.g. [0057] and FIG. 6B, 6C, 7A, 7B, 8A and 8B).

15. As to claim 11, Talagala discloses wherein storing the first data block and the second data block comprises using a storage pool allocator ("Storage Controller 401" – e.g. FIG. 2 and 3, [0033], [0059]).
16. As to claim 12, it is rejected using the same rationale as for the rejection of claim 1.
17. As to claim 13, it is rejected using the same rationale as for the rejection of claim 6.
18. As to claim 14, Talagala discloses wherein the indirect block is stored using a storage pool allocator ("Storage Controller 401" –e.g. FIG. 2 and 3, [0033]).
19. As to claim 15, Talagala discloses a method for retrieving data in a data block (abstract), comprising: obtaining an indirect block comprising a first block pointer comprising a first stored checksum and a first data block location and a second block pointer comprising a second stored checksum and a second data block location ([0055], [0058], [0059] and FIG. 6B, 6C, 7A, 7B, 8A and 8B); obtaining the first data block using

Art Unit: 2135

the first data block location ([0017], [0059]); calculating the checksum for the first data block to obtain a calculated checksum ([0059], FIG. 6C, 7B and 8B); retrieving the data from the first data block, if the stored checksum equals the calculated checksum ("The block received from the READ may be compared to the checksum received from the READ " –e.g., [0040], [0059], FIG. 2, 3, 6C, 7B and 8B); and performing an appropriate action, if the first stored checksum is not equal to the calculated checksum ("The block received from the READ may be compared to the checksum received from the READ (e.g. by recalculating the checksum from the data and comparing the new checksum to the read checksum). If they do not match, an error may be detected immediately" –e.g., see, [0040], [0059], FIG. 2, 3, 6C, 7B and 8B).

20. As to claim 16, Talagala discloses wherein the calculated checksum is calculated using a storage pool allocator ("Storage Controller 401" –e.g. FIG. 2, 3 and 6C, [0033], [0059]).

21. As to claim 17, Talagala discloses a method for storing and retrieving data blocks (abstract), comprising: storing a first data block and a second data block ([0017], [0059]); obtaining a first data block location and a second data block location ([0017], [0059]); calculating a first data block checksum for the first data block and a second data block checksum for the second data block ([0058]); storing a first block pointer in an indirect block, wherein the first block pointer comprises the first data block location and the first data block checksum ([0059], FIG. 6C, 7B and 8B); storing a second block

pointer in the indirect block, wherein the second block pointer comprises the second data block location and the second block checksum ([0059], FIG. 6C, 7B and 8B); obtaining the indirect block comprising the first block pointer and the second block pointer ([0059] and FIG. 6C, see also [0013], [0052]-[0053], [0060], [0065] and FIG. 6A, 6B, 6C); obtaining the first data block using the first data block location stored in the first block pointer ([0059] and FIG. 6C, see also [0013], [0052]-[0053], [0060], [0065] and FIG. 6A, 6B, 6C); calculating the checksum for the first data block to obtain a calculated checksum ([0016], [0017], [0058] – [0059]); retrieving data from the first data block, if the first data block checksum stored in the first block pointer equals the calculated checksum ([0040], [0059], FIG. 2, 3, 6C, 7B and 8B); and performing an appropriate action, if the first data block checksum is not equal to the first calculated checksum ([0040], [0059], FIG. 2, 3, 6C, 7B and 8B).

22. As to claim 18, Talagala discloses a system for storing data blocks, comprising: a storage pool comprising a first data block, a second data block and a first indirect block referencing the first data block and the second data block, wherein the first indirect block comprises a first data block checksum and a first data block location stored in a first block pointer, and a second data block checksum and a second data block location stored in a second block pointer ([0016], [0017], [0058] – [0059]); and a storage pool allocator configured to store the first data block, the second data block and the first indirect block in the storage pool (FIG. 2 and 3, [0033], "Storage Controller 401").

Art Unit: 2135

23. As to claim 19, Talagala discloses the system further comprising: a second indirect block, comprising a first indirect data block checksum and a first indirect block location, wherein the storage pool allocator is further configured to store the second indirect block in the storage pool (Talagala teaches writing a block having a checksum and a location in PGT (Parity Group Table) which have a "segment" field to store a location and a "checksum" field to store a checksum and provides which contain parity/stripe group tables having entries for multiple blocks of data stored in a storage pool, -e.g. [0059] and FIG. 6C, 7B and 8B).

24. As to claim 20, Talagala discloses further comprising: a data management unit configured to assemble the first indirect block and request the storage pool allocator to store the first indirect block ("Storage Controller 401" –e.g. FIG. 2 and 3, [0033]).

25. As to claim 21, Talagala discloses wherein the storage pool comprises at least one storage device ("array of storage devices 410" –e.g. FIG. 2 and 3, [0033]).

26. As to claim 22, it is rejected using the same rationale as for the rejection of claim 8.

27. As to claim 23, it is rejected using the same rationale as for the rejection of claim 9.

28. As to claim 24, it is rejected using the same rationale as for the rejection of claim 10.

29. As to claim 25, it is rejected using the same rationale as for the rejection of claim 1.

30. As to claim 26, Talagala discloses a network system having a plurality of nodes, comprising: a storage pool comprising a first data block, a second data block and a first indirect block referencing the first data block and the second data block, wherein the first indirect block comprises a first data block checksum and a first data block location stored in a first block pointer, and a second data block checksum and a second data block location stored in a second block pointer ([0016], [0017], [0058] – [0059]); and a storage pool allocator configured to store the first data block, the second data block and the first indirect block in the storage pool (FIG. 2 and 3, [0033], “Storage Controller 401”), wherein the storage pool is located on any one of the plurality of nodes ([0017], [0059], [0051]), and wherein the storage pool allocator is located on any one of the plurality of nodes ([0017], [0033], [0059], [0051]) and FIG. 6C).

Response to Arguments

31. Applicant’s arguments filed 01 August 2007, have been fully considered but they are not persuasive.

Art Unit: 2135

32. Regarding the applicant's remarks that "Talagala fails to disclose the use of a **hierarchical tree structure** in which an indirect block references multiple data blocks" as recited in the independent claims; it is the Examiner's position that this limitation is not found within the scope of the claim language as the independent claims recite "a first indirect block referencing the first data block and the second data block" (Talagala discloses the limitation as "an indirection map (e.g., block remapping table) matches virtual block address to physical block address. Block-level checksums may be provided in the indirection map" –e.g. [0059] and FIG. 6C; "each valid PGT entry also includes a back pointer to the next entry in a parity group so that the first physical segment in a parity group is linked to the second physical segment in that parity group, and the second physical segment to the third and so on, until the last physical segment contains the parity data for that parity group. The physical segment that contains the parity data is linked back to the first physical segment in the parity group, thereby creating a circular list for that parity group" –e.g. [0056], Applicant should note that a circular linked list will comprise "a first physical segment in a parity group" which comprises a first indirect block referencing the first and second data blocks and so on as claimed by applicant)

33. Regarding the applicant's remarks that: "Talagala fails to disclose the use of a metaslab in a file system"; it is the Examiner's position that Talagala discloses this limitation as (As defined by Applicant, metaslabs are "contiguous regions of data" in which "the storage space in the storage pool is divided" (Specification, [0032]);

Art Unit: 2135

therefore, Applicant should note that these metaslabs may comprise any amount of contiguous data, such as segments or blocks into which a storage pool is divided, as described by Talagala. Note that Talagala discloses "stripes of data" within an array storage devices (Figure 4 and [0043] and explains that a "stripe" of data is analogous to a "parity group" [0063], lines 10-11, wherein configuration information for each of the "parity groups" is stored in a "PGT (parity group table)" (FIG. 6C, 7B and 8B) and further explains "the indirection map may also include a parity group pointer for each data block that points to a next member of that parity group" –e.g. [0013], herein "when a READ or • WRITE command is received for a block(s), the appropriate PGT entry is accessed locate the blocks in the disk drives" –e.g. [0058]).

34. Examiner's note: Examiner has cited particular columns and line numbers in the references as applied to the claims above for the convenience of the applicant. Although the specified citations are representative of the teachings in the art and are applied to the specific limitations within the individual claim, other passages and figures may be applied as well. It is respectfully requested from the applicant, in preparing the responses, to fully consider the references in entirety as potentially teaching all or part of the claimed invention as well as the context of the passage as taught by the prior art or disclosed by the examiner.

Conclusion

35. **THIS ACTION IS MADE FINAL.** Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

36. Any inquiry concerning this communication or earlier communications from the examiner should be directed to Suman Debnath whose telephone number is 571 270 1256. The examiner can normally be reached on 8 am to 5 pm.


If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Kim Y. Vu can be reached on 571 272-3859. The fax phone number for the organization where this application or proceeding is assigned is 571-273-8300.

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only.

Art Unit: 2135

For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative or access to the automated information system, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

69
SD



KIM VU
PATENT EXAMINER
ART CENTER 2100